

Improved methods for detecting gravitational waves associated with short gamma-ray bursts

A. R. Williamson,¹ C. Biwer,² S. Fairhurst,¹ I. W. Harry,^{1,2} E. Macdonald,¹ D. Macleod,^{1,3} and V. Predoi¹

¹*School of Physics and Astronomy, Cardiff University, Cardiff, UK*

²*Department of Physics, Syracuse University, Syracuse, NY 13244, USA*

³*Louisiana State University, Baton Rouge, LA 70803, USA*

(Dated: October 23, 2014)

In the era of second generation ground-based gravitational wave detectors, short gamma-ray bursts (GRBs) will be among the most promising astrophysical events for joint electromagnetic and gravitational wave observation. A targeted search for gravitational wave compact binary merger signals in coincidence with short GRBs was developed and used to analyze data from the first generation LIGO and Virgo instruments. In this paper, we present improvements to this search that enhance our ability to detect gravitational wave counterparts to short GRBs. Specifically, we introduce an improved method for estimating the gravitational wave background to obtain the event significance required to make detections; implement a method of tiling extended sky regions, as required when searching for signals associated to poorly localized GRBs from *Fermi* Gamma-ray Burst Monitor or the InterPlanetary Network; and incorporate astrophysical knowledge about the beaming of GRB emission to restrict the search parameter space. We describe the implementation of these enhancements and demonstrate how they improve the ability to observe binary merger gravitational wave signals associated with short GRBs.

I. INTRODUCTION

Gamma-ray bursts (GRBs) are amongst the most energetic electromagnetic events in the universe, observed isotropically across the sky and up to cosmological redshifts [1]. An apparent bimodality observed in the duration and spectral hardness of GRBs – long-soft and short-hard — suggests more than one class of progenitors [2]. The mergers of compact binary systems composed of two neutron stars or a neutron star and a black hole have long been proposed as possible progenitors of short GRBs [3, 4]. Short GRB variability timescales are small, indicating a compact source [5], while the observed offsets from their host galaxies agree with that expected for a population of compact binary mergers and not of core-collapse supernovae [6], the widely accepted progenitors of most long GRBs. The recent detection of a kilonova associated with GRB130603B [7, 8] has further supported the compact merger hypothesis. For an in-depth review of short GRB science, see eg. [9, 10].

If short GRBs are indeed compact binary mergers, they are a very interesting class of events for gravitational wave (GW) astronomy, since such compact binary mergers are also strong emitters of GWs [11, 12]. GW observations of GRBs will make possible direct observation of the central engines that power these events, a feat that electromagnetic observations alone cannot achieve due to circumburst material and ejecta [13]. The observation of a short GRB provides the time and sky position of a potential GW source. A targeted search for a binary merger GW signal, informed by the GRB observation, need only search a small fraction of the parameter space of an un-triggered, full-sky, binary merger search. Consequently, it is possible to significantly reduce the detection threshold for the targeted GRB search [14], thereby increasing the sensitivity of the search.

An analysis pipeline has been developed specifically for performing the targeted search for binary mergers associated with GRBs [15]. At its heart, this is a matched-filtering analysis [16] that makes use of the well understood gravitational waveforms emitted during binary merger to search for a sig-

nal in data from the operational GW detectors. The analysis makes use of the known sky location of the GRB and the relative GW detector sensitivities to appropriately time shift and weight the data streams from the individual detectors to perform a *coherent* analysis. This is in contrast to a *coincidence* analysis, which individually analyzes the data from the different detectors and then performs a coincidence check (see e.g. [17]). By coherently combining the data, it is possible to isolate data streams containing the two gravitational wave polarizations. In the case where data from more than two GW detectors are used, the other, orthogonal contributions provide a *null stream* which helps to reduce the impact of non-stationary noise. The analysis pipeline also makes use of signal consistency tests to check that anything causing a large signal-to-noise ratio (SNR) is consistent with a putative GW signal. As well as using these tests to reject noise transients, or *glitches*, they are also used to down-weight events which are more consistent with a noise transient than a signal. The significance of the results is computed by comparing the result for the six seconds around the time of the GRB with data from surrounding times, thereby calculating the probability of obtaining an event with a specific SNR due to noise alone.

The analysis pipeline described above has been used to carry out numerous GRB searches on data from the initial Laser Interferometer Gravitational Wave Observatory (LIGO) [18] and Virgo [19] detectors. Searches associated with short GRBs observed by the *Swift* [20] and *Fermi* [21] satellites have been performed [22, 23], as well as those observed by the InterPlanetary Network (IPN) [24, 25]. None of these analyses made a GW detection in conjunction with an observed short GRB, which was not surprising given the sensitivity of the initial detectors — tens of Mpc for binary merger signals — and the typical distances to GRBs — a median redshift of 0.7 and a closest measured redshift of 0.1, implying a distance of 500 Mpc. The second generation of GW detectors, Advanced LIGO [26] and Advanced Virgo [27], are due to begin observing in 2015 with sensitivities that will increase, over time, to approximately ten times greater than

those of the first generation detectors [28]. The prospects for detection of GW signals associated with GRBs with the advanced detectors are promising [14, 29–31].

With the realistic prospect of a joint GW–GRB observation in the coming years, we have made a number of changes and improvements to the analysis pipeline. These enhancements are critical to optimizing the potential for observing binary merger GW signals in coincidence with electromagnetically observed short GRBs. Most importantly, we have improved the ability of the pipeline to estimate the significance of rare events. To this end, we have introduced the ability to perform time-shifted analyses, whereby the data from the different detectors are shifted by several seconds relative to each other and the analysis is repeated. This allows us to measure the background of the search to lower than one part in 10^5 , a level that would be required for an unambiguous detection claim [32].

The targeted search introduced in [15] makes use of the sky location of the source, but places no restrictions on the orientation of the binary. Observations of short GRBs indicate that the gamma-ray jet is beamed, with most observations favoring a beaming angle of 30° or smaller [33], with the GRB jet emitted orthogonal to the orbital plane of the binary. Thus, it is natural to incorporate this into the search by restricting the search to binaries which are observed to be (close to) face-on. This restriction reduces the parameter space of the search, providing an increase in sensitivity.

While *Swift* provides typical GRB localizations with arcminute accuracy [20], *Fermi*’s Gamma-ray Burst Monitor (GBM) and the IPN often provide localizations to significantly larger regions of the sky [21, 24]. For these GRBs, the 3σ confidence sky localization region can be tens of square degrees so it no longer suffices to search a single sky point. We describe a method of searching over a grid of sky points that cover the uncertainty region, and demonstrate that by using this grid we can efficiently search for a signal originating from any point in the sky patch. This method has already been used in the searches described in [23, 25].

The layout of this paper is as follows. In Section II we review the coherent analysis pipeline introduced in [15]. In Section III we describe and demonstrate the improved background calculation. In Section IV we describe how searching only for signals with a narrow opening angle can improve search sensitivity. In Section V we describe how to tile search points on the sky and explore how this improves sensitivity for GRBs with larger sky localization regions. Finally, in Section VI we discuss the implications of our results.

II. PIPELINE SUMMARY

The targeted, coherent search was described in detail in [15]. Here, we provide a brief review of the analysis pipeline in order to provide the necessary background for the following sections. Where possible, we follow the same notation as the original paper. For a more detailed description see [15] itself.

The targeted, coherent search is carried out whenever an observed GRB is detected during a time that at least two grav-

itational wave detectors are operating and have good quality data for a sufficiently long period of time either side of the GRB. In practice, we search for a signal in a 6 second window covering 5 seconds before to 1 second after the Earth crossing time of the GRB called the *on-source* window. However, we require additional data around this time in order to perform the analysis to ensure that the detectors were operating stably at the time of the GRB and to provide a good estimate of the detector sensitivity at the time. Our ability to detect a GW signal associated with a GRB depends upon both the stationary noise background and also the non-stationary noise transients in the data which might mask a signal. The data surrounding the on-source time is used to evaluate both of these.

A. Multi-Detector Matched Filter

The pipeline performs a modelled gravitational wave search for compact binary inspiral signals. To this end, a bank of template waveforms [34, 35] that densely cover the mass parameter space is used to perform a matched-filter analysis [16]. The targeted GRB search performs a coherent analysis whereby the data streams from different detectors are combined while performing the matched-filtering and a network SNR is calculated directly. The majority of searches for binary merger GWs perform the matched-filter independently on individual interferometer data streams before comparing the resulting *triggers* to search for coincident events (see e.g. [17]). The coherent analysis affords several benefits. First, by performing the analysis coherently, we combine the detector data to produce two data streams which are sensitive to the two gravitational wave polarizations. Any other, orthogonal data streams will necessarily contain only noise and can either be ignored, or used to eliminate noise transients which will often contribute power to these *null streams*. Additionally, by combining the data from the detectors at the time of analysis, we will accumulate power from all detectors, not just those which produced a trigger above threshold. It was shown in [15] that a coherent analysis provides an improvement in sensitivity over the coincident one, but the search is more computationally costly than a coincident one. A targeted GRB search, where both the sky location and arrival time of the signal are constrained is ideal for performing the more sensitive, coherent analysis.

The amplitude of a GW signal from a non-precessing binary may be decomposed into two polarizations, denoted $+$ and \times , as

$$h_+(t) = \mathcal{A}^1 h_0(t) + \mathcal{A}^3 h_{\pi/2}(t), \quad (1a)$$

$$h_\times(t) = \mathcal{A}^2 h_0(t) + \mathcal{A}^4 h_{\pi/2}(t). \quad (1b)$$

Here, h_0 and $h_{\pi/2}$ denote the two phases of the waveform, which depend upon the binary masses as well as the coalescence time of the signal. These are calculated using the post-Newtonian formalism [34]. In the analysis presented here, we restrict to non-spinning components. However, the search is easily extended to binaries with spins aligned with the orbital angular momentum by simply generating additional templates

to cover the spin parameter space (see e.g. [36, 37]). The amplitude terms for an inspiral GW signal are

$$\mathcal{A}^1 = \frac{D_0}{D} \frac{(1 + \cos^2 \iota)}{2} \cos 2\phi_0 \cos 2\psi - \frac{D_0}{D} \cos \iota \sin 2\phi_0 \sin 2\psi, \quad (2a)$$

$$\mathcal{A}^2 = \frac{D_0}{D} \frac{(1 + \cos^2 \iota)}{2} \cos 2\phi_0 \sin 2\psi + \frac{D_0}{D} \cos \iota \sin 2\phi_0 \cos 2\psi, \quad (2b)$$

$$\mathcal{A}^3 = -\frac{D_0}{D} \frac{(1 + \cos^2 \iota)}{2} \sin 2\phi_0 \cos 2\psi - \frac{D_0}{D} \cos \iota \cos 2\phi_0 \sin 2\psi, \quad (2c)$$

$$\mathcal{A}^4 = -\frac{D_0}{D} \frac{(1 + \cos^2 \iota)}{2} \sin 2\phi_0 \sin 2\psi + \frac{D_0}{D} \cos \iota \cos 2\phi_0 \cos 2\psi. \quad (2d)$$

These terms are dependent on 4 variables: the source distance, D ; the coalescence phase, ϕ_0 ; the polarization angle, ψ ; and the inclination angle, ι . D_0 is a scaling distance (usually 1 Mpc). It is worth noting that, for any set of amplitudes \mathcal{A}^μ , there is a unique set of $\{D, \iota, \phi_0, \psi\}$, up to reflection and rotation symmetry.

The GW signal seen by a detector X is a combination of the two polarizations, each weighted by an antenna power pattern factor $F_{\{+, \times\}}$ [38], which describes the relative response of the detector to each polarization,

$$h^X(t) = F_+^X h_+(t^X) + F_\times^X h_\times(t^X). \quad (3)$$

Here, t^X is the time of arrival of the signal at detector X , which will depend upon a fiducial arrival time (for example at the geocenter) and the relative location of the detector and source.

In matched-filtering analysis the inner products between a template gravitational waveform time series $h(t)$ and detector data stream time series $s(t)$ are calculated. In general, the inner product between two such time series, a^X and b^X , is given by

$$(a^X | b^X) = 4 \operatorname{Re} \int_0^\infty \frac{\tilde{a}^X(f) \cdot \tilde{b}^X(f)^*}{S_h^X(f)}, \quad (4)$$

where $S_h^X(f)$ is the noise power spectral density in detector X , and $\tilde{a}(f)$ denotes the Fourier transform of the time series $a(t)$. For binary merger signals, the two phases h_0 and $h_{\pi/2}$ are orthogonal, in the sense that

$$(h_0 | h_{\pi/2}) = 0. \quad (5)$$

For a network of detectors, we define the multi-detector inner product as the sum of the single detector inner products,

$$(\mathbf{a} | \mathbf{b}) \equiv \sum_{X=1}^d (a^X | b^X), \quad (6)$$

where d denotes the number of detectors in the network. The multi-detector log-likelihood is then defined as,

$$\ln \Lambda = (\mathbf{s} | \mathbf{h}) - \frac{1}{2} (\mathbf{h} | \mathbf{h}) = \left[\mathcal{A}^\mu (\mathbf{s} | \mathbf{h}_\mu) - \frac{1}{2} \mathcal{A}^\mu \mathcal{M}_{\mu\nu} \mathcal{A}^\nu \right], \quad (7)$$

where $\mathbf{h} = (\mathbf{F}_+ \mathbf{h}_0, \mathbf{F}_\times \mathbf{h}_0, \mathbf{F}_+ \mathbf{h}_{\pi/2}, \mathbf{F}_\times \mathbf{h}_{\pi/2})$, and the matrix

$$\mathcal{M}_{\mu\nu} \equiv (\mathbf{h}_\mu | \mathbf{h}_\nu). \quad (8)$$

Maximizing this likelihood ratio over the amplitude parameters \mathcal{A}_μ , we obtain the maximized coherent SNR,

$$\rho_{\text{coh}}^2 \equiv 2 \ln \Lambda|_{\text{max}} = [(\mathbf{s} | \mathbf{h}_\mu) \mathcal{M}^{\mu\nu} (\mathbf{s} | \mathbf{h}_\nu)], \quad (9)$$

where $\mathcal{M}^{\mu\nu}$ is the inverse of the matrix $\mathcal{M}_{\mu\nu}$.

The coherent SNR forms the basis of the detection statistic and has a χ^2 background distribution with 4 degrees of freedom. The four degrees of freedom correspond to the four components of the gravitational wave signal – the 0 and $\pi/2$ phases of the two polarizations. This becomes more transparent if we work in the dominant polarization frame. In this frame, the network is maximally sensitive to the + polarization and the two polarizations are orthogonal. Then, the coherent SNR can be re-expressed as

$$\rho_{\text{coh}}^2 = \frac{(\mathbf{s} | \mathbf{F}_+ \mathbf{h}_0)^2 + (\mathbf{s} | \mathbf{F}_\times \mathbf{h}_{\pi/2})^2}{(\mathbf{F}_+ \mathbf{h}_0 | \mathbf{F}_+ \mathbf{h}_0)} + \frac{(\mathbf{s} | \mathbf{F}_\times \mathbf{h}_0)^2 + (\mathbf{s} | \mathbf{F}_+ \mathbf{h}_{\pi/2})^2}{(\mathbf{F}_\times \mathbf{h}_0 | \mathbf{F}_\times \mathbf{h}_0)}. \quad (10)$$

In Gaussian noise, the coherent SNR would be the detection statistic. Events with a larger coherent SNR would be less likely to be due to noise fluctuations and consequently more likely to be due to a GW signal. However, in real data GW signals are not the only cause of deviations from the background distribution. Noise transients, or *glitches*, also contribute to the background. Although glitches will not typically mimic template waveforms, if they are large enough they will still produce a large SNR. Consequently, we must use a number of consistency tests to eliminate or down-weight triggers that are unlikely to be due to a GW signal incident upon the detector network.

B. Signal Consistency

Matched filtering alone leads to the identification of a large number of triggers, many of which are purely due to non-Gaussian noise transients present in the data streams. Such noise transients may be discarded by performing signal consistency tests across the individual detectors that make up the network. Here, we briefly describe the different tests used in the analysis.

1. Null Stream Consistency

Null stream consistency makes use of one or more null data streams or, in the case of this pipeline, the related null SNR

statistic. This is simply the SNR observed in the detector network that is *not* consistent with the signal model:

$$\rho_N^2 \equiv \sum_X \rho_X^2 - \rho_{\text{coh}}^2, \quad (11)$$

where ρ_X is the SNR in detector X . For a signal which matches the template waveform, there will be no signal power in the null SNR, and it will be χ^2 distributed with $2d - 4$ degrees of freedom due to the presence of noise. An incoherent, non-Gaussian transient noise event will contribute to the null SNR and consequently a large null SNR is used to eliminate spurious events via a hard cut if

$$\begin{aligned} \rho_{\text{null}} &> 5.25, & \rho_{\text{coh}} &\leq 20 \\ \rho_{\text{null}} &> \frac{\rho_{\text{coh}}}{5} + 5.25, & \rho_{\text{coh}} &> 20 \end{aligned} \quad (12)$$

2. Single detector thresholds

Noise transients are, by their nature, events which occur in a single detector. Conversely, gravitational wave events will lead to signal power being distributed among all detectors in the network. We can use this difference to further reduce the background due to glitches. The most effective, and most straightforward, method is simply to require that a signal is observed with an SNR above threshold (typically four) in at least two detectors. This serves to eliminate the majority of glitches, which have power in only one detector, with very little effect on signals.

3. χ^2 Tests

When matched-filtering identifies a trigger with a large SNR there is necessarily some component of the data which matches the signal $h(t)$. If the trigger is caused by a noise glitch, there is likely to be an additional, orthogonal component of the data which is not well described by Gaussian noise. χ^2 tests are designed to eliminate glitch triggers by identifying power that is not consistent with either signal or Gaussian noise. To do so, we introduce a set of basis waveforms T^i which are orthonormal and also orthogonal to the signal waveform $h(t)$. Specifically, we require

$$(\mathbf{T}_\mu^i | \mathbf{T}_\nu^j) = \delta^{ij} \delta_{\mu\nu} \quad \text{and} \quad (\mathbf{T}_\mu^i | \mathbf{h}_\nu) = 0, \quad (13)$$

where μ, ν refer to the waveform components and i, j the waveforms that comprise the basis for the χ^2 test. We then construct a χ^2 statistic as

$$\chi^2 = \sum_{\mu=1}^4 \sum_{i=1}^N (\mathbf{T}_\mu^i | \mathbf{s})^2. \quad (14)$$

In the presence of a signal that matches the template waveform (or no signal), the statistic will be χ^2 distributed with $4N$ degrees of freedom. If the data contains some additional, non-Gaussian noise the χ^2 value will be elevated provided that the

set of templates T^i captures at least a fraction of the power contained in the glitch. Triggers with a large χ^2 value are discarded. In practice it is far from trivial to choose the set of waveforms T^i so that they are both orthonormal and orthogonal to $h(t)$, and match a variety of non-Gaussianities. Three different χ^2 tests have been implemented in the analysis:

- i *Frequency bins*: The test waveforms T^i are generated by chopping up the template $h(t)$ into $(N+1)$ sub-templates in the frequency domain, each of which contains an equal amount of power. From these, we generate N orthonormal waveforms which are also orthogonal to $h(t)$.
- ii *Template bank*: The test waveforms T^i are taken from the template-bank of binary merger waveforms used in the search. In general, these will not be orthogonal to $h(t)$, but it is straightforward to subtract the part proportional to $h(t)$. However, it is more difficult to render the waveforms T^i orthonormal. In practice we do not attempt to do so, but instead use an empirical threshold based on an effective number of degrees of freedom.
- iii *Autocorrelation*: The test waveforms T^i are simply copies of the waveform $h(t)$ offset in time from the original. As with the template bank, it is straightforward to remove the component of T^i that is proportional to $h(t)$. We do not attempt to orthonormalize the T^i and again empirically set the threshold.

4. Re-weighted SNR

In addition to discarding triggers which fail the signal consistency test described above, we also re-weight the SNR of triggers based on the values of the χ^2 tests and null SNR. This allows us to better differentiate signals from noise background. The re-weighting is chosen such that the SNR of signals will be unaffected while those noise triggers which do not match well with the template waveform will be down-weighted. We perform two sets of down-weighting. Firstly, with the frequency bin χ^2 values,

$$\rho_{\chi^2} = \begin{cases} \rho_{\text{coh}} & \chi^2 \leq n_{\text{dof}} \\ \frac{\rho_{\text{coh}}}{\left\{ \left[1 + \left(\frac{\chi^2}{n_{\text{dof}}} \right)^3 \right] / 2 \right\}^{1/6}} & \chi^2 > n_{\text{dof}} \end{cases}, \quad (15)$$

and then the null SNR,

$$\rho_{\text{rw}} = \begin{cases} \rho_{\chi^2} & \rho_{\text{null}} \leq 4.25 \\ \frac{\rho_{\chi^2}}{\rho_{\text{null}} - 3.25} & \rho_{\text{null}} > 4.25 \end{cases}. \quad (16)$$

This re-weighted SNR value is the detection statistic used for evaluating candidate events.

We note that the χ^2 re-weighted SNR given in Eq. (15) is different from the one used in the original paper [15]. In particular, the exponents in the denominator have been changed. In the process of developing an all-sky, all-time coherent analysis [39], it was found that the original re-weighting left a small tail of high SNR noise events. These had not been

observed in the GRB search previously, due to the limited amount of data used in the analyses. By using a re-weighted SNR identical to the one used in the all-sky coincidence search [17], we were able to eliminate the high SNR events. The same re-weighting has now been applied in the GRB search.

C. Event significance

The analysis described above is performed for all template waveforms in the template bank covering the mass space. For each template, the re-weighted SNR is calculated. The template producing the largest re-weighted SNR during the on-source window is retained as the event candidate.

The significance of this event is calculated using the data before and after the time of the GRB, which is designated *off-source*. This data will not contain a signal corresponding to the GRB and is also unlikely to contain a GW signal from the same sky position which is unassociated with the GRB, thus any events occurring in the off-source will be due to background noise. In a typical search we use approximately an hour of data for the off-source, and split this into trials with durations equal to that of the on-source window. This gives us a means of characterizing the background noise in our detector network around the time of the GRB. The significance of the on-source event is determined by calculating the false alarm probability, or p-value. This is simply the fraction of off-source trials with an event of equal or greater significance than in the on-source.

To evaluate the sensitivity of the pipeline to finding GW signals in the data around the time of the GRB, we inject a number of simulated signals into the off-source data. The simulated signals are drawn randomly from an astrophysically motivated distribution of distances, component masses and spins and binary inclination. The simulated signals are compact binary merger waveforms at 3.5 post-Newtonian order [34, 40], where one component of the binary is taken to be a neutron star and the second either a neutron star or black hole. The efficiency of the analysis at recovering these signals provides a measure of pipeline performance and produces a lower limit on the distance to which the pipeline is sensitive.

D. Example GRB

In the remainder of this paper, we will illustrate the various pipeline developments using example analyses based upon the GRB 100928A, which was observed by the *Swift* Burst Alert Telescope (BAT) [41, 42]. No other *Swift* instrument observed this GRB as the spacecraft was unable to slew to the sky position of the prompt burst due to a Sun observing constraint. It was not detected by *Fermi* or any other gamma ray sensitive instrument.

We have chosen this GRB for a number of reasons. Virgo and both LIGO detectors were operational and had ample science quality data either side of the GRB time. Specifically, 5264 seconds of coherent network data between 01:34:35 and

03:02:19 UTC on 28 September 2010 was available for analysis purposes. Additionally, the BAT localized the burst to a point on the sky (RA = 223.037°, Dec = −28.542°) where both LIGO detectors were approximately equally sensitive, and where Virgo had good sensitivity. Furthermore, this position was known accurately, with a 90% confidence radius of only 2.3 arcminutes.

III. BACKGROUND ESTIMATION

To make a confident detection statement, we must establish that the probability of an observed event being due to noise alone is very small. This requires a detailed understanding of the search background generated by both Gaussian detector noise and non-stationary transients. We do this by looking at the data around the time of the GRB. We make the reasonable assumption that the off-source data contains no GW signal originating from the same location on the sky and has, on average, the same statistical properties as the detector network background during the on-source period. Thus, the off-source data provides a means of characterizing the background noise in the detector network at the time of the GRB.

The false alarm probability (FAP) associated to the on-source event, with re-weighted SNR ρ^* , is the probability of having a more significant event in any randomly chosen 6 seconds of data. This is calculated by counting the fraction of background trials which have an event with $\rho > \rho^*$,

$$\text{FAP} = \frac{N(\rho > \rho^*)}{N_{\text{BG}}}, \quad (17)$$

where N_{BG} denotes the total number of background trials. In the standard approach, we simply split the background into as many 6 second trials as possible, so the number of background trials is given by $N_{\text{BG}} = T_{\text{off}}/T_{\text{on}}$.

The standard analysis makes use of approximately an hour of data around the time of the GRB, leading to a lower limit on the FAP of around 10^{-3} . For the majority of GRBs, this will be sufficient to demonstrate that there is no candidate GW event associated to a particular GRB. However, when there is an interesting candidate, a FAP of 10^{-3} is not sufficient to warrant a detection claim, and further background trials are required to more accurately evaluate the significance.

What would be an acceptable FAP to support a detection claim? In particle physics, the standard level is a “5-sigma” observation, or 1 in 3 million. In a recent gravitational wave search [32], a simulated signal was added to the data and recovered with a false alarm rate of 1 in 7,000 years, which was deemed sufficient to claim evidence for a detection. Translating this to a GRB search equates to a FAP of 3×10^{-6} for one of the 50 short GRBs observed each year. Alternatively, we might consider the chance of there being an observable signal around the time of a GRB. In [14], this was estimated to be around 1% for the advanced detector network operating at design. Clearly, a detection candidate would require a FAP much lower than the probability of observing a signal. All arguments point to requiring a minimum of 10^5 background trials to assess the significance of a detection candidate.




FIG. 1. FAP as a function of the re-weighted SNR detection statistic for a search performed for GRB 100928A, using time slides to reach $\text{FAP} < 10^{-5}$. The figure shows the background estimated with off-source only (787 trials) plotted in orange Y; the short slide analysis (8917 trials) plotted in green x; both long and short slides (267185 trials) plotted in blue +. With short slides alone, we can estimate a significance of 1 part in 10^4 while long and short slides give a background estimate to 1 in 4×10^6 . The shaded regions show the 95% Jeffreys credible interval for each case, which assumes each time slide is a statistically independent trial. For clarity of presentation we have only plotted the 20 loudest trials for each search.

To reach a significance level of 10^{-5} , we require further background trials. The most straightforward approach would be to simply extend the off-source analysis to incorporate one week of data. While in principle this is possible, the typical duration of continuous operation for the detectors is on the order of hours. Furthermore, the data quality is known to change between different stretches of data, so a week of off-source data may not accurately characterize the data at the time of the GRB. In addition, extending the off-source data to one week would increase the computational cost of the analysis by a factor of several hundred, rendering it impractical to estimate the background promptly. Consequently, an alternative method is required. To obtain an improved estimate of the network background, we instead artificially time shift the data from the different detectors and repeat the analysis. These time shifts are always significantly longer than the light travel time between detectors and the signal auto-correlation time and typical glitch durations (all well under one second), so that GW signals will not appear coherently in the time-shifted analysis.

We are able to increase the number of background trials performed by an order of magnitude, with minimal impact on the computational cost, thereby allowing us to estimate FAPs to around 10^{-4} . This is achieved by time shifting the SNR time series of the individual detectors prior to performing the coherent analysis. In the analysis, the detector data is split into sections, typically of 128 seconds length, which are match filtered to produce a (complex) SNR time series for each detector. These are then combined according to Eq. (9) to calcu-

late the coherent SNR time series. A *short slide* is performed by introducing relative time-shift between the detectors' SNR time series prior to computing the coherent SNR. For the example GRB, we leave the H1 data alone, shift the L1 data by multiples of 6 seconds and the V1 data by multiples of 12 seconds. This allows for ten time shifted analyses to be performed. Since calculating the single detector SNR time series is the most computationally costly part of the analysis, short slides have a relatively small computational cost. In Fig. 1, we show the improvement in background estimation afforded by the inclusion of the short slides.

We have also implemented *long slides* which involve permuting the data segments prior to analysis. Unfortunately, this does require repeating the analysis, so the computational cost increases linearly with the number of long slides. However, it is possible to perform short slides within each long slide. Thus, we only require around ten long slides in order to achieve a background estimate of 10^{-5} .

Fig. 1 shows FAP as a function of re-weighted SNR for the analysis of GRB 100928A. This shows that any on-source event with $\rho_{\text{rw}} > 8.5$ would have a FAP at the 10^{-5} level. We have, however, assumed that all time slides are independent. In reality, all time slides are formed from different combinations of the same detector data streams, and so are not statistically independent at all. A more rigorous treatment of FAP uncertainty when dealing with time slides would likely show far larger 95% credible intervals for all cases, however it is not clear how to implement such a treatment for this search [43].

It is interesting to compare the background for the GRB search with the all-sky coincidence search [32]. This will allow us to estimate the sensitivity improvement offered by the targeted, coherent search. For the all sky search, the background is one event per year at an SNR of 10 decreasing by two orders of magnitude per unit increase in SNR.¹ Interestingly, the background for the targeted, coherent search, as shown in Fig. 1, falls off at the same rate. In both cases, this is significantly slower than expected in Gaussian noise, suggesting that both pipelines are affected in a similar way by the non-Gaussian transients in the data. The background for the all-sky coincidence search translates to a FAP of 10^{-3} in six seconds of data at an SNR of 8.2. In comparison, the targeted, coherent search achieves this background at an SNR of 7.3. While both of these are re-weighted SNR measurements, and the details of the pipelines differ, the analysis methods have much in common, so it is reasonable to compare the results. Thus, the coherent analysis provides approximately a 13% reduction in the SNR at a given FAP.

We can use this to estimate the benefit of performing the GRB search. To do so, we compare against a simple analysis that just examines the results of the all-sky search for triggers within the 6 second on-source window. The comparison of FAPs above shows that the targeted, coherent search would

¹ This is taken from Figure 3 in [32], which shows a background of around 0.2 events per year at SNR of 10. However, we must also apply a trials factor of six, as described in the paper, to give a background of 1 event per year at this SNR.




FIG. 2. Comparison between the $+$, \times amplitude terms as a function of inclination angle ι . Note that even at 30° the difference is only $\sim 1\%$.

identify a candidate event with a 13% lower SNR, or equivalently at a 13% greater distance. In addition, the targeted, coherent search applies lower single detector SNR thresholds of 4, rather than 5.5, and it includes the SNR contribution from all detectors, even if they did not produce a trigger above threshold. For the case of GRB 100928A, a signal near the detection threshold would be unlikely to register as a trigger in the Virgo detector, and the coherent analysis would register about 10% greater SNR by incorporating the power from Virgo. This implies that the targeted, coherent search provides approximately a 23% increase in distance sensitivity over a search that simply looks for a coincident GW trigger from the all sky search. This equates to approximately a doubling of the event rate.

IV. SEARCH FOR SIGNALS WITH NARROW OPENING ANGLES

Short GRBs are believed to be beamed phenomena [44, 45], with prompt γ -ray emission concentrated along collimated jets normal to the orbital plane. These jets are expected to have opening angles of $< 30^\circ$ [33]. Therefore, it may be reasonable to assume that observed short GRB progenitor systems have their orbital angular momenta nearly parallel with the line-of-sight, corresponding to system orbital inclinations $\iota \sim 0$ or $\iota \sim \pi$ with respect to the observer.

In Eq. (2), we see that the GW amplitudes depend linearly on $\cos \iota$ and $(1 + \cos^2 \iota)/2$. For a binary inclination close to $\iota = 0$, both of these tend towards unity. In Fig. 2, we plot both amplitude factors as a function of ι . This serves to highlight the fact that the amplitudes vary almost identically with ι , up to an angle of 30° , by which time they differ by only $\sim 1\%$. Even at 45° , the two amplitudes differ by only 6%. Consequently for GRB signals, it is reasonable to treat the amplitude factors as equal and to approximate the signal as left circularly polarized. Similarly, when $\iota \sim 180^\circ$, the two terms agree up

to an overall sign and the signal is right circularly polarized.

It is therefore convenient to introduce a single amplitude and phase to describe the signal as

$$\tilde{D} = \frac{D}{\cos \iota} \quad \text{and} \quad \chi_{l,r} = \phi_0 \pm \psi. \quad (18)$$

Then, for $\iota \approx 0$, the amplitudes simplify to

$$\mathcal{A}^1 \approx \mathcal{A}^4 \approx -\frac{D_0}{\tilde{D}} \cos 2\chi_l \equiv \mathcal{B}_1, \quad (19a)$$

$$\mathcal{A}^2 \approx -\mathcal{A}^3 \approx \frac{D_0}{\tilde{D}} \sin 2\chi_l \equiv \mathcal{B}_2, \quad (19b)$$

and similar for $\iota \approx 180^\circ$. As expected, the circularly polarized GW signal is then dependent upon two amplitudes \mathcal{B}_1 and \mathcal{B}_2 (or, equivalently, a single overall amplitude and phase),

$$h_+(t) = \mathcal{B}_1 h_0(t) - \mathcal{B}_2 h_{\pi/2}(t), \quad (20a)$$

$$h_\times(t) = \mathcal{B}_2 h_0(t) + \mathcal{B}_1 h_{\pi/2}(t). \quad (20b)$$

rather than the original four amplitudes \mathcal{A}^μ .

Substituting these expressions into Eq. (7), and working in the dominant polarization, we obtain,

$$\ln \Lambda = \mathcal{B}_1 (\mathbf{s} | \mathbf{F}_+ \mathbf{h}_0 + \mathbf{F}_\times \mathbf{h}_{\pi/2}) + \mathcal{B}_2 (\mathbf{s} | \mathbf{F}_\times \mathbf{h}_0 + \mathbf{F}_+ \mathbf{h}_{\pi/2}) - \frac{1}{2} [\mathcal{B}_1^2 + \mathcal{B}_2^2] [(\mathbf{F}_+ \mathbf{h}_0 | \mathbf{F}_+ \mathbf{h}_0) + (\mathbf{F}_\times \mathbf{h}_0 | \mathbf{F}_\times \mathbf{h}_0)] \quad (21)$$

It is straightforward to maximize over the amplitude parameters $\mathcal{B}_{1,2}$ to obtain

$$\rho_{\text{coh}}^2 = \frac{\alpha^2 + \beta^2}{(\mathbf{F}_+ \mathbf{h}_0 | \mathbf{F}_+ \mathbf{h}_0) + (\mathbf{F}_\times \mathbf{h}_0 | \mathbf{F}_\times \mathbf{h}_0)}, \quad (22)$$

where

$$\alpha = (\mathbf{s} | \mathbf{F}_+ \mathbf{h}_0) + (\mathbf{s} | \mathbf{F}_\times \mathbf{h}_{\pi/2}), \quad (23a)$$

$$\beta = (\mathbf{s} | \mathbf{F}_\times \mathbf{h}_0) - (\mathbf{s} | \mathbf{F}_+ \mathbf{h}_{\pi/2}). \quad (23b)$$

The calculation proceeds in an analogous manner for $\iota \sim 180^\circ$, with the signal now right, rather than left, polarized. After maximization, the coherent SNR takes the same form as Eq. (22), but with

$$\alpha = (\mathbf{s} | \mathbf{F}_+ \mathbf{h}_0) - (\mathbf{s} | \mathbf{F}_\times \mathbf{h}_{\pi/2}), \quad (24a)$$

$$\beta = (\mathbf{s} | \mathbf{F}_\times \mathbf{h}_0) + (\mathbf{s} | \mathbf{F}_+ \mathbf{h}_{\pi/2}). \quad (24b)$$

The motivation for performing the search for only circularly polarized waveforms is to further reduce the noise background and thereby increase the sensitivity of the search. Additionally, restricting to circularly polarized waveforms provides us with an additional *null stream* that can be used to reject noise glitches. Prior to assessing the improvement in real data, it is useful to evaluate the expected benefit in Gaussian noise. The original search has four free amplitude parameters \mathcal{A}^μ , and the coherent SNR in the absence of a signal is χ^2 distributed with four degrees of freedom. When restricting to circular polarization, there are two free parameters \mathcal{B}_μ and the coherent SNR in Gaussian noise will be χ^2 distributed with two degrees




FIG. 3. The background significance against detection statistic for a search performed for GRB 100928A. In red \times , we plot the background calculated using the circular polarization restriction and in blue $+$ we plot the background from the un-restricted search. In both cases, we perform time shifts of the data as discussed in Section III. Over a broad range of SNR values, the circular polarization restriction reduces the background by a factor of three. Equivalently, the required SNR to achieve a given FAP is reduced by about 0.25, equating to a 3% increase in the distance sensitivity of the search. For clarity of presentation we have only plotted the loudest 50 trials for each search.

of freedom. However, we must now search over both left and right circularly polarized signals, which leads to a doubling of the number of trials.² Comparison of these distributions, for a large number of trials, suggests restricting to circular polarization should result in at a decrease in FAP of around one order of magnitude at fixed SNR, or an increase in sensitivity at fixed FAP of roughly 5%.

In Fig. 3 we plot the FAP as a function of SNR for the circularly polarized and un-restricted searches. Over a broad range of SNRs we observe a reduction in the background of a factor of three, corresponding to an increase in sensitivity of around 3% at a given FAP. This improvement is less significant than might have been expected in Gaussian data, and only serves to emphasize that the data we are using is not Gaussian, and events in the tail of the distribution are due to noise transients in the data.

Interestingly, we have noticed that the most significant background triggers in the circular search do not correspond to outliers in the un-restricted search. This is likely due to how the pipeline selects triggers. It first applies a clustering method to choose the trigger with the largest coherent SNR in a given time window, before applying signal consistency tests to the trigger which may lead to it being discarded or the SNR

re-weighted. Consequently, it is possible that loud events in the un-restricted search do not survive in the circular analysis, and vice versa.

We have demonstrated that restricting to circularly polarized signals can provide a small improvement in the search sensitivity and, furthermore, that it is a reasonable approximation given our current understanding of GRB beaming. We note that a 3% improvement in distance reach corresponds to a 10% increase in the rate of observable signals.

V. SEARCHING A PATCH OF THE SKY

Short GRBs are localized to sky error boxes of varying sizes by different satellites. This has implications for the targeted GW search following up on these events. For example, the BAT instrument aboard NASA's *Swift* satellite is capable of localizing to 1-4 arcminutes [20], while the typical GW localization region is several square degrees or larger [28, 46]. Thus, we may follow up a BAT trigger by searching only a single point on the sky since the GRB localization is significantly better than the sky resolution of the GW search. However, the GBM aboard NASA's *Fermi* satellite often localizes GRBs to far larger patches of the sky [21]. The 3σ confidence regions are roughly circular, with a radius of several degrees. Additionally, the IPN localizes GRBs by triangulation with a number of satellites [24]. Depending upon the number of satellites observing the event and their relative positions, the localizations can range from under a square degree to hundreds or even thousands of square degrees. For poorly localized short GRBs observed by *Fermi* or IPN, the GRB localization will be comparable to, or larger than, the typical GW localization region. Consequently, it is no longer appropriate to treat the GRB localization as a single point in the sky, and we must extend the GW search to cover the entire confidence region.

The targeted, coherent GW search makes use of the sky location in two ways. Firstly, and most importantly, it is the sky location which determines the relative arrival time of a signal at the detectors in the network. These time delays are used to appropriately shift the data prior to coherently combining them in the search. Using the incorrect sky location will cause the signals from different detectors to be mis-aligned in time. Secondly, the detector sensitivities, encoded in the antenna response factors $F_{(+,\times)}$, depend upon the location of the source relative to the detector. The use of incorrect $F_{(+,\times)}$ will lead to the wrong weighting of detector data streams in the coherent SNR and signal power being present in the null stream.

We can estimate when the single sky point search will not be sufficient. To do so, let us consider only the loss in SNR arising from timing offsets. In a matched-filter search, the recovered SNR in a detector falls off as

$$\rho(dt)^2 \approx \rho_o^2 [1 - (2\pi\sigma_f)^2 dt^2], \quad (25)$$

where σ_f is the signal bandwidth, which is typically around 100 Hz for a binary merger signal [46]. Thus, a timing offset of $\delta t = 0.5\text{ms}$ will lead to a 5% loss in SNR in a single detector.

² The left and right circular waveforms are only orthogonal when the network is equally sensitive to both polarizations. For most sky locations, this is not the case, and the two trials are *not* independent leading to a further reduction in the expected background.

Given a network of N detectors, $D_{\{1,\dots,N\}}$, let \mathbf{r}_i denote the location of the detector and t_i be the arrival time of the GW signal at detector i from a GRB at the central location of the sky patch. The distance between two detectors is

$$d_{ij} = \|\mathbf{r}_j - \mathbf{r}_i\|, \quad (26)$$

and the light travel time between them is

$$T_{ij} = d_{ij}/c. \quad (27)$$

The difference in the arrival time of the signal at two detectors, τ_{ij} , is calculated as [47],

$$\tau_{ij} = t_i - t_j = \frac{1}{c} (\mathbf{r}_i - \mathbf{r}_j) \cdot \mathbf{w} \equiv T_{ij} \cos \alpha, \quad (28)$$

where \mathbf{w} is the unit wave vector describing the direction of propagation of the source, and α is the angle between the line connecting the detectors and the direction to the source.

It is then straightforward to calculate the change in time delay with a change in the angle α as

$$\delta\tau_{ij} = \sqrt{T_{ij}^2 - \tau_{ij}^2} \delta\alpha. \quad (29)$$

So, for a source lying on the line connecting the two detectors, the time delay τ_{ij} between detectors is maximal and changes only quadratically with the change in the location of the source. In contrast, for a source which lies on the zero time delay plane, $\tau_{ij} = 0$, a change in location will induce the largest time offset.

Once we select the maximum time offset δt that we are willing to tolerate, it is straightforward to calculate the required angular spacing of the sky points as

$$\delta\alpha = \min_{i,j} \left[\frac{2\delta t}{\sqrt{T_{ij}^2 - \tau_{ij}^2}} \right]. \quad (30)$$

Here, the factor of two arises because δt is the largest single detector time offset. We typically choose $\delta t = 0.5\text{ms}$. The two LIGO detectors are separated by a light travel time of 10ms, while LIGO and Virgo are separated by around 25ms, which sets the angular scale to around 2° for the LIGO detectors and 1° between LIGO and Virgo. In practice, the resolution is usually determined by the detector pair (D_i, D_j) for which the GRB target location has smallest relative arrival time difference.

The circular grid is generated by placing rings of points spaced by $\delta\alpha$, starting at the center, with the final ring passing the 3σ confidence radius. An example of such a grid is shown in Fig. 4 (full grid). Each ring will have $2\pi n/\delta\alpha$ points, where $n = 0$ labels the central point and increases as we move outwards. The method of covering the patch is based upon the one introduced in [48]. In the analysis, each point in the grid is treated independently, with the single-detector data streams time shifted appropriately for the given sky location. The coherent SNR and signal consistency tests are calculated with the appropriate detector responses, F_+ and F_\times , for that




FIG. 4. An example patch of sky points projected onto the celestial sphere. The blue filled circles show the full grid, while the empty circles are those few points that map to unique differences in signal arrival time between LIGO's Hanford and Livingston detectors. The parsed points do not form a straight line, but this is simply due to an artifact of the parsing routine and has no effect on the grid reduction.

sky point. As with the background estimation, searching over points in the sky patch is performed *after* the computationally dominant step of calculating the single detector SNR time series. Consequently, GRBs observed by *Fermi* GBM, requiring around hundred sky points, are processed in approximately double the time required for the *Swift* GRBs with a single sky point.

To demonstrate the efficacy of searching over a sky patch, we repeated the analysis of GRB 100928A, but used a typical *Fermi* GBM 3σ localization uncertainty radius of 15° [21]. The sky patch for the search contained 178 search points in total. When performing simulations, the location of each source was chosen randomly from a normal distribution with width 5° , ie. $\sim 99\%$ of simulated signals were within the 15° radius 3σ localization region. In the rest of this section, we use the results for binary neutron star (BNS) systems exclusively to illustrate the effect of a search over a patch of the sky. However, similar effects are observed in searches for neutron star-black hole (NSBH) systems.

In Fig. 5, we show the search efficiency as a function of distance for three different searches: a single point search with simulations spread over the 0.036° *Swift* BAT sky patch; a single point search with simulations spread over a typical 15° *Fermi* GBM sky patch; and a grid of points covering the GBM sky patch with simulations spread over the patch. In all cases the efficiency is calculated at the SNR of the loudest background event in the *short slide* analysis. If we perform the search using only a point at the center of the *Fermi* localization region, the results are poor: across the whole range of distances, the search efficiency is never greater than 40%, even for nearby signals which have large SNRs. The reason for this lies in the signal consistency tests discussed in Section II. At the incorrect sky location, the signal does not match the




FIG. 5. The fraction of artificially injected binary neutron star signals found louder than the loudest background event as a function of injected distance. The three curves represent three observational scenarios for a three detector network comprised of Virgo and both LIGO interferometers. In the scenario mimicking a BAT GRB (black solid line, error radius = 0.036°) the pipeline searches a single point on the sky and finds 90% of signals within 20 Mpc. In the two scenarios mimicking a GBM GRB we see that by searching over a patch of points covering the large error box of 15° radius (red dashed line) the pipeline performs nearly as well as for the BAT GRB for signals below 15 Mpc. This is in stark contrast to the previous treatment for GBM-like GRBs (blue dotted line), which searched a single point at the center of the error box resulting in very poor rates of injection recovery. The increased number of trials resulting from multiple sky points leads to a tail of background events louder than any seen of the BAT single point search, reducing the overall sensitivity of the patch search.

template due to inevitable time offsets between them and the signal will be recovered with a different phase in each of the detectors. Consequently, the coherent SNR will not correctly reflect the total signal power and this will lead to increased values of the signal consistency tests, and in particular lead to a significant amount of power in the null SNR. Thus, many signals at smaller distances are found with large total SNR values but also have sufficient power in the null SNR statistic to fail the hard cut in Eq. (12) and are therefore rejected as potential detections.

The sensitivity of the search over the *Fermi* error region is almost the same as the search over just the *Swift* point at small distances, but decreases more rapidly for quieter signals at larger distances. For example, the distance at which we achieve 50% efficiency is reduced by 10%. This loss in sensitivity can be attributed to the fact that we have searched over numerous sky points thereby increasing the number of background trials. In this example, we obtain a loudest background event with re-weighted SNR value of 8.33 compared to 7.51 for a single point, which is consistent with the loss of reach of the search.

This method of placing a grid of points in the sky has already been used in the analysis of *Fermi*-detected GRBs dur-




FIG. 6. The fraction of artificially injected binary neutron star signals found louder than the loudest background event using only the LIGO Hanford and Livingston detectors, plotted as a function of injected distance. As in Fig. 5, we plot a scenario mimicking a BAT GRB (black solid line, error radius = 0.036°) where the pipeline searches a single point on the sky. In this case, the pipeline finds 90% of signals within 18 Mpc. In the scenario where a GBM GRB with error box of 15° radius is searched at a single point (blue dotted line), we see poor signal recovery performance at small distances due to signal consistency effects, similar to the three detector case. The difference between the full patch of search points (red dashed line) and a set of points covering unique time delays between sites (green dot-dashed line) is noticeable at small distances, with the use of incorrect antenna response factors causing a drop in performance for the parsed patch. Again, the increased number of trials resulting from multiple sky points leads to a tail of background events louder than any seen of the BAT single point search, reducing the overall sensitivity of multiple point searches.

ing LIGO Science Run 6 and Virgo Science Runs 2 & 3. An analogous method was used to perform the search over the irregular sky patches produced by the IPN [49].

A. Two-site time-delay degeneracy

In the case of a two-site detector network, for example the LIGO-only network, the ability to resolve independent sky locations is vastly reduced. With a single baseline between sites, multiple sky locations will map to the same difference in signal arrival time. Thus, when moving across the sky patch, there will be one direction where only the antenna response factors $F_{(+,\times)}$ change, and not the time delays, while in the orthogonal direction both will change. With two detectors, after maximising over the \mathcal{A}^μ , the values of $F_{(+,\times)}$ drop out of the coherent SNR expression. This is not immediately obvious, but can be understood by noting that for a two detector search, there are four degrees of freedom in both the coincident and coherent searches. Therefore, *any* observed amplitude and phase in the two detectors is consistent with an astrophysical signal; there is no null stream. Then, the size of the sky grids

can be significantly reduced, to represent only those sky locations that map to unique time-delays between observatory sites. Fig. 4 shows an example result of parsing the circular sky maps to remove degeneracies in time-delay. For the map shown, only 20% of the points are required to uniquely span the allowed time-delays between the LIGO sites, allowing a reduction in cost in the analysis for two-site GRB analyses.

Unfortunately, once we restrict to circularly polarized signals, as described in Section IV, the restriction to a single time-delay line is no longer appropriate. Now, there are only two free signal amplitudes, which cannot match arbitrary amplitude and phase measurements in the two detectors. Thus the detector response functions are again enter into the construction of the coherent SNR and the circular null stream.

In Fig. 6, we show the sensitivity of the search performed using only the two LIGO detectors in Hanford and Livingston and incorporating an inclination restriction. As before, we plot the *Swift* search results – where both the simulated signals and search are restricted to a single sky point – as a reference. Next we consider the GRB localized to a typical *Fermi* GBM error region. When searching over the full *Fermi* sky patch, there is again a degradation of the sensitivity due to a tail of loud background events (a maximum SNR of 8.12 compared to 7.25 for the single point search). However, searching of a single sky point leads to a dramatic loss of sensitivity, with only 60% of nearby signals being recovered. By searching over only the one dimensional time-delay space, we recover the majority of this sensitivity, but do observe a small drop in efficiency at low distances due to the use of incorrect antenna response factors.

VI. DISCUSSION

The advanced LIGO and Virgo detectors will be sensitive to binary merger signals from hundreds of Mpc, or even Gpc in the case of NSBH systems. These distances are then comparable to those of the closest GRBs and a joint observation in the coming years is a distinct possibility. In this paper, we have presented details of an improved GW–GRB search that implements several new features that will facilitate joint observations. The work extends that in [15] in three distinct ways. First, we have introduced a method of time-shifting the background data in order to estimate false alarm probabilities lower than 10^{-5} . An event of this significance, or greater, will likely be required to claim the first joint GW–GRB observation. Critically, we have seen that there is no “tail” of rare, high SNR events that would hinder a detection claim. Second, we have developed a method for searching over extended regions of the sky, rather than just a single point. The majority of short GRB observations are currently made with the *Fermi* GBM detector, which typically localizes events to tens of square degrees. With the capability of searching sky patches, we can now achieve a comparable sensitivity for *Fermi* GBM bursts as to those which are localized to arc-second accuracy with *Swift*. Third, we have made use of astrophysical priors on GRB jets to restrict the search to nearly face-on binaries whose gravitational wave signal will be circularly polarized.

We have shown that this provides a small, but significant, improvement in sensitivity of the search. Taken together, these improvements to the search mean it is ready to be deployed in the future on advanced GW detector data at the time of short GRBs. Nonetheless, there are several additional features that we plan to implement in the near future, which we describe below.

We would like to provide rapid, Gamma-ray Coordination Network (GCN) style alerts of the GW search results for GRBs. For these to be useful, the analysis must be completed as rapidly as possible. This can be achieved by a simple re-ordering of the analysis to prioritize the on-source analysis with short time slides only to provide an initial result, with a FAP measured to 10^{-4} within an hour or two of the GRB. Subsequent improvements on FAP measurement and search sensitivity using simulated signals will follow later if a promising GW candidate is found. To make this feasible in the long term, a process is being developed to automatically launch the analysis upon receipt of a GCN alert.

The current search makes use of template waveforms appropriate for binaries with non-spinning components. For neutron stars, this is a reasonable approximation as they are expected to have low spins which will not greatly affect the waveform [37]. However, in an NSBH system, the black hole spin can have a significant effect on the emitted waveform. The component of the spin aligned with the orbital angular momentum will affect the rate at which the binary inspirals [50], while the orthogonal spin components will lead to precession of the system [51]. It has been shown [36, 52] that using waveforms which incorporate the effects of aligned spins can greatly enhance the sensitivity of a search to NSBH systems. Furthermore, when the spin is aligned with the orbital angular momentum, the waveforms simplify to the form given in Eq. (1). Thus, it is straightforward to simply extend the template bank to include these waveforms and incorporate the effects of aligned spins. Including the aligned spin contribution to the waveform will aid the sensitivity of the search. It is not so straightforward to incorporate precession effects. Thankfully, precession typically has a less significant effect on the waveform when the binary is observed close to face on [53], so that will reduce the importance of precession for the GRB search. Nonetheless, we would like to incorporate these effects. In [54], we investigated a method of extending the search to waveforms with precession. In the future, we will identify the regions of parameter space where the spin-aligned waveforms do not provide good sensitivity to precessing signals and complete the development of the analysis in [54] to provide a sensitive search over these parts of the parameter space.

Not all NSBH mergers are expected to emit electromagnetically [55–57]. Under most scenarios, electromagnetic emission requires the formation of a torus around the central black hole. In [58], the region of black hole masses and spins which might give rise to this torus was investigated. Around half of the NSBH parameter space will not lead to torus formation, under any reasonable model of a neutron star equation of state, and can therefore be eliminated from the analysis. By eliminating these templates from the analysis, we can reduce

both the computational cost and search background. This is a feature we plan to implement in the near future.

The advanced detectors are sensitive to signals from 10 Hz upwards [26], in comparison to 40 Hz for the initial detectors. A BNS system will take 1,000 seconds to evolve from 10 Hz to merger. Consequently, the search must be extended to deal with longer duration templates, in order to capture all of the available power in the signal. This can be achieved by extending the lengths of the analyzed data segments in the search, although significant changes will be required to handle 1,000 second templates. However, the early advanced detector runs are not expected to obtain the full low-frequency sensitivity [28]. For these, a search beginning at 25 Hz (templates of 90 second duration) would capture the vast majority of the signal power. This can be achieved relatively easily, and will be available for the initial runs that are expected in 2015.

ACKNOWLEDGMENTS

The authors would like to thank the LIGO Scientific Collaboration and the Virgo Collaboration for allowing access to the data used in generating the results in Sections III to V. In addition, they thank Patrick Sutton for recommending 100928A as a suitable example GRB, and Dipongkar Talukder for comments on the manuscript. This work was supported by the UK Science and Technologies Funding Council through grant ST/L000962/1, the Royal Society and the US National Science Foundation through grants PHY-0847611, PHY-1205835, PHY-1404395, and AST-1333142.

-
- [1] W. S. Paciesas, C. A. Meegan, G. N. Pendleton, M. S. Briggs, C. Kouveliotou, *et al.*, *Astrophys.J.Suppl.* **122**, 465 (1999), arXiv:astro-ph/9903205 [astro-ph].
 - [2] C. Kouveliotou, C. A. Meegan, G. J. Fishman, N. P. Bhyat, M. S. Briggs, *et al.*, *Astrophys.J.* **413**, L101 (1993).
 - [3] D. Eichler, M. Livio, T. Piran, and D. N. Schramm, *Nature* **340**, 126 (1989).
 - [4] R. Narayan, B. Paczynski, and T. Piran, *Astrophys.J.* **395**, L83 (1992), arXiv:astro-ph/9204001 [astro-ph].
 - [5] M. J. Rees and P. Meszaros, *Astrophys.J.* **430**, L93 (1994), arXiv:astro-ph/9404038 [astro-ph].
 - [6] R. P. Church, A. J. Levan, M. B. Davies, and N. Tanvir, *Mon.Not.R.Astron.Soc.* **413**, 2004 (2011), arXiv:1101.1088 [astro-ph.HE].
 - [7] E. Berger, W. Fong, and R. Chornock, *Astrophys.J.* **774**, L23 (2013), arXiv:1306.3960 [astro-ph.HE].
 - [8] N. R. Tanvir, A. J. Levan, A. S. Fruchter, J. Hjorth, K. Wiersema, *et al.*, *Nature* **500**, 547 (2013), arXiv:1306.4971 [astro-ph.HE].
 - [9] E. Nakar, *Phys.Rept.* **442**, 166 (2007), arXiv:astro-ph/0701748 [astro-ph].
 - [10] E. Berger, *Ann. Rev. Astron. Astrophys.* **52**, 43 (2014), arXiv:1311.2603 [astro-ph.HE].
 - [11] C. Cutler, T. A. Apostolatos, L. Bildsten, L. S. Finn, E. E. Flanagan, *et al.*, *Phys.Rev.Lett.* **70**, 2984 (1993), arXiv:astro-ph/9208005 [astro-ph].
 - [12] M. Shibata and K. Taniguchi, *Phys.Rev.* **D77**, 084015 (2008), arXiv:0711.1410 [gr-qc].
 - [13] I. Bartos, P. Brady, and S. Marka, *Class.Quant.Grav.* **30**, 123001 (2013), arXiv:1212.2289 [astro-ph.CO].
 - [14] J. Clark, H. Evans, S. Fairhurst, I. W. Harry, E. Macdonald, D. Macleod, P. J. Sutton, and A. R. Williamson, (2014), arxiv:1409.8149 [astro-ph.HE].
 - [15] I. W. Harry and S. Fairhurst, *Phys.Rev.* **D83**, 084002 (2011), arXiv:1012.4939 [gr-qc].
 - [16] L. A. Wainstein and V. D. Zubakov, *Extraction of Signals from Noise* (Prentice-Hall, Englewood Cliffs, NJ, 1962).
 - [17] S. Babak, R. Biswas, P. R. Brady, D. Brown, K. Cannon, *et al.*, *Phys.Rev.* **D87**, 024033 (2013), arXiv:1208.3491 [gr-qc].
 - [18] B. P. Abbott *et al.* (LIGO Scientific Collaboration), *Rept.Prog.Phys.* **72**, 076901 (2009), arXiv:0711.3041 [gr-qc].
 - [19] T. Accadia *et al.* (VIRGO Collaboration), *JINST* **7**, P03012 (2012).
 - [20] S. D. Barthelmy, L. M. Barbier, J. Cummings, E. E. Fenimore, N. Gehrels, *et al.*, *Space Sci.Rev.* **120**, 143 (2005), arXiv:astro-ph/0507410 [astro-ph].
 - [21] C. Meegan, G. Lichti, P. Bhat, E. Bissaldi, M. S. Briggs, *et al.*, *Astrophys.J.* **702**, 791 (2009), arXiv:0908.0450 [astro-ph.IM].
 - [22] J. Abadie *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Astrophys.J.* **715**, 1453 (2010), arXiv:1001.0165 [astro-ph.HE].
 - [23] J. Abadie *et al.* (LIGO Scientific Collaboration), *Astrophys.J.* **760**, 12 (2012), arXiv:1205.2216 [astro-ph.HE].
 - [24] K. Hurley, T. Cline, I. Mitrofanov, E. Mazets, S. Golenetskii, *et al.*, *AIP Conf.Proc.* **662**, 473 (2003), arXiv:astro-ph/0201374 [astro-ph].
 - [25] J. Aasi *et al.* (The LIGO Scientific Collaboration, the Virgo Collaboration), *Phys.Rev.Lett.* **113**, 011102 (2014), arXiv:1403.6639 [astro-ph.HE].
 - [26] G. M. Harry (LIGO Scientific Collaboration), *Class.Quant.Grav.* **27**, 084006 (2010).
 - [27] F. Acernese *et al.*, (2009), <https://pub3.ego-gw.it/itf/tds/file.php?callFile=VIR-0027A-09.pdf>.
 - [28] J. Aasi *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), (2013), arXiv:1304.0670 [gr-qc].
 - [29] J. Abadie *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Class.Quant.Grav.* **27**, 173001 (2010), arXiv:1003.2480 [astro-ph.HE].
 - [30] A. Dietz, N. Fotopoulos, L. Singer, and C. Cutler, *Phys.Rev.* **D87**, 064033 (2013), arXiv:1210.3095 [gr-qc].
 - [31] H.-Y. Chen and D. E. Holz, *Phys.Rev.Lett.* **111**, 181101 (2013), arXiv:1206.0703 [astro-ph.CO].
 - [32] J. Abadie *et al.* (LIGO Collaboration, Virgo Collaboration), *Phys.Rev.* **D85**, 082002 (2012), arXiv:1111.7314 [gr-qc].
 - [33] A. N. Guelbenzu, S. Klose, J. Greiner, D. A. Kann, T. Kruehler, *et al.*, *Astron.Astrophys.* **548**, A101 (2012), arXiv:1206.1806 [astro-ph.HE].
 - [34] L. Blanchet, *Living Rev.Rel.* **17**, 2 (2014), arXiv:1310.1528 [gr-qc].
 - [35] B. J. Owen and B. Sathyaprakash, *Phys.Rev.* **D60**, 022002 (1999), arXiv:gr-qc/9808076 [gr-qc].
 - [36] I. W. Harry, A. H. Nitz, D. A. Brown, A. Lundgren, E. Ochsner, *et al.*, *Phys.Rev.* **D89**, 024010 (2014), arXiv:1307.3562 [gr-qc].

- [37] D. A. Brown, I. Harry, A. Lundgren, and A. H. Nitz, *Phys.Rev.* **D86**, 084017 (2012), arXiv:1207.6406 [gr-qc].
- [38] P. Jaranowski, A. Krolak, and B. F. Schutz, *Phys.Rev.* **D58**, 063001 (1998), arXiv:gr-qc/9804014 [gr-qc].
- [39] D. Macleod *et al.*, In preparation (2014).
- [40] L. Blanchet, T. Damour, G. Esposito-Farese, and B. R. Iyer, *Phys.Rev.Lett.* **93**, 091101 (2004), arXiv:gr-qc/0406012 [gr-qc].
- [41] V. D’Elia, S. T. Holland, E. A. Hoversten, H. A. Krimm, C. B. Markwardt, M. H. Siegel, M. Stamatikos, G. Stratta, and E. Troja, *GRB Coordinates Network* **11310**, 1 (2010).
- [42] H. A. Krimm, S. D. Barthelmy, W. H. Baumgartner, J. R. Cummings, V. D’Elia, E. E. Fenimore, N. Gehrels, C. B. Markwardt, D. M. Palmer, T. Sakamoto, G. Sato, M. Stamatikos, J. Tueller, and T. N. Ukwatta, *GRB Coordinates Network* **11312**, 1 (2010).
- [43] M. Was, M.-A. Bizouard, V. Brisson, F. Cavalier, M. Davier, *et al.*, *Class.Quant.Grav.* **27**, 015005 (2010), arXiv:0906.2120 [gr-qc].
- [44] W. Fong, E. Berger, B. D. Metzger, R. Margutti, R. Chornock, *et al.*, *Astrophys.J.* **780**, 118 (2014), arXiv:1309.7479 [astro-ph.HE].
- [45] A. Panaitescu, *Mon.Not.R.Astron.Soc.Lett.* **367**, L42 (2006), arXiv:astro-ph/0511588 [astro-ph].
- [46] S. Fairhurst, *New J.Phys.* **11**, 123006 (2009), arXiv:0908.2356 [gr-qc].
- [47] O. Rabaste, E. Chassande-Mottin, and A. Pai, (2009), arXiv:0905.4832 [gr-qc].
- [48] J. Aasi *et al.* (LIGO Scientific Collaboration, VIRGO Collaboration), *Phys.Rev.* **D89**, 122004 (2014), arXiv:1405.1053 [astro-ph.HE].
- [49] V. Predoi and K. Hurley (IPN Collaboration, LIGO Scientific Collaboration, Virgo Collaboration), *J.Phys.Conf.Ser.* **363**, 012034 (2012), arXiv:1112.1637 [gr-qc].
- [50] P. Schmidt, M. Hannam, and S. Husa, *Phys.Rev.* **D86**, 104063 (2012), arXiv:1207.3088 [gr-qc].
- [51] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, *Phys.Rev.* **D49**, 6274 (1994).
- [52] P. Ajith, *Phys.Rev.* **D84**, 084037 (2011), arXiv:1107.1267 [gr-qc].
- [53] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, *et al.*, *Phys.Rev.Lett.* **113**, 151101 (2014), arXiv:1308.3271 [gr-qc].
- [54] I. W. Harry and S. Fairhurst, *Class.Quant.Grav.* **28**, 134008 (2011), arXiv:1101.1459 [gr-qc].
- [55] F. Foucart, *Phys.Rev.* **D86**, 124007 (2012), arXiv:1207.6304 [astro-ph.HE].
- [56] N. Stone, A. Loeb, and E. Berger, *Phys.Rev.* **D87**, 084053 (2013), arXiv:1209.4097 [astro-ph.HE].
- [57] C. Palenzuela, L. Lehner, M. Ponce, S. L. Liebling, M. Anderson, *et al.*, *Phys.Rev.Lett.* **111**, 061105 (2013), arXiv:1301.7074 [gr-qc].
- [58] F. Pannarale and F. Ohme, *Astrophys.J.* **791**, L7 (2014), arXiv:1406.6057 [gr-qc].